

**IN THE UNITED STATES
PATENT AND TRADEMARK OFFICE**

PATENT APPLICATION

JC564 U.S. PTO
09/22/99

APPLICANT(S): Henry Esmond Butterworth
Carlos Francisco Fuente
Robert Frank Maddock

TITLE: Data Processing Systems and
Method for Processing Work
Items in Such Systems

**ASSISTANT COMMISSIONER FOR PATENTS
WASHINGTON, D.C. 20231**

SIR:

Enclosed are the following papers relating to the above-named application for patent:

Application - 16 pages of Specification, 4 pages of Claims and 1 page of Abstract
3 Sheets of informal drawing(s)

Associate Power of Attorney

1 Assignment with Cover Sheet

Declaration and Power of Attorney

Information Disclosure Statement with Form PTO-1449 and cited reference(s)

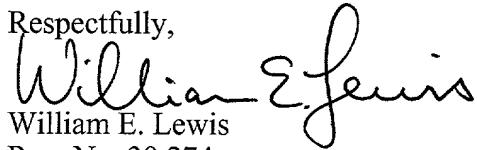
A Certified Copy of Foreign Application for which priority is claimed (UK Patent Application No. 9906501.3, filed March 23, 1999)

Check in the amount of \$800.00 to cover the filing fee and Assignment recordation fee

CLAIMS AS FILED				
	NO. FILED	NO. EXTRA	RATE	CALCULATIONS
Total Claims	11-20 =	0	x \$18 =	\$0
Independent Claims	3-3 =	0	x \$78 =	\$0
Multiple Dependent Claim(s), if applicable			\$260 =	\$0
Basic Fee				\$760
			TOTAL FEE:	\$760

Please file the application and please find enclosed a check in the amount of \$800 (which includes \$760 to cover the filing fee and \$40 to cover the Assignment recordation fee). In the event of non-payment or improper payment of a required fee, the Commissioner is authorized to charge or to credit **Ryan & Mason, L.L.P. Deposit Account No. 50-0762** as required to correct the error. Duplicate copies of this letter are enclosed.

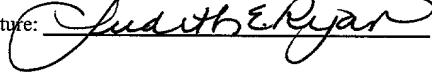
Please address all correspondence to: **William E. Lewis, Ryan & Mason, L.L.P., 90 Forest Avenue, Locust Valley, New York 11560.** Telephone calls should be made to the under-signed attorney at (516) 759-2946.

Respectfully,

 William E. Lewis
 Reg. No. 39,274

Date: September 22, 1999
Ryan & Mason, L.L.P.
90 Forest Avenue
Locust Valley, New York 11560

"Express Mail" Label No. EJ937152848US
 I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated below and is addressed to: Assistant Commissioner for Patents, Washington, D.C. 20231.

Date of Deposit: September 22, 1999

Signature: 

JC564 U.S. PTO
09/22/99

**DATA PROCESSING SYSTEMS AND METHOD FOR PROCESSING
WORK ITEMS IN SUCH SYSTEMS**

Field of the Invention

5 The present invention relates generally to processing of work items in data processing systems and more particularly to the scheduling of tasks to process such work items.

10 Background of the Invention

In modern data processing systems there are two well established methods for the communication required between hardware and software for software control of hardware devices within the system: interrupts and polling. With interrupts, each hardware device signals that there is work for the software to do by asserting an interrupt line which causes the software flow-of-control to be diverted to an interrupt handler which handles the interrupt. An interrupt is typically handled by masking (or disabling) the interrupt and scheduling a task for later execution which will service the requesting device. Once the device has been serviced, the interrupt is unmasked (or enabled) again ready to allow the generation of further interrupts. With polling, the device driver continually checks the status of the device to determine whether it needs to be serviced. If so, the device is serviced.

15
20
25

25

With polling based systems, there is a trade-off between polling frequently so as to service the device with low latency and polling infrequently so as to keep the overhead of polling to a minimum. With interrupt-based systems, the device is serviced with low latency and low overhead when it is relatively idle but when the device becomes heavily utilised and is generating many interrupts the increased interrupt overhead can make the interrupt method inefficient. Generally speaking therefore, an interrupt based system is more efficient at relatively low device utilisation and a polling based system is more efficient at relatively high device utilisation. Neither technique is especially well adapted to the situation where the device presents varying amounts of work to the system.

Attempts have been made in the art to reduce the interrupt overhead in interrupt-based systems. In IBM Technical Disclosure Bulletin Vol 38 No 7 pp 497-500, a technique is described that is designed to reduce the number of interrupts generated by a LAN adapter. One type of interrupt (transmit complete) which does not impact response time is turned off thus reducing system disruption and CPU utilisation. An alternative method for identifying and reporting transmit completions is provided. Although the described technique goes some way to reducing interrupt overhead in the described environment, it necessitates the provision of an alternative method to handle the disabled interrupt.

It would be desirable to have a system having low latency at low workload levels and low overhead at high workload levels.

5

Disclosure of the Invention

According to a first aspect of the invention there is provided a method for processing work items in a data processing system comprising: generating an interrupt in response to receipt of a work item in the system; servicing the generated interrupt to schedule a task for later processing of the work item, without re-enabling the interrupt; subsequently executing the task to process the work item; and speculatively scheduling a further task for processing of any work items that are subsequently received in the system.

The present invention is thus contrasted with prior systems in which, when an interrupt is serviced and a task scheduled for later execution, the interrupt is enabled/unmasked. Thus in prior systems, the receipt of further work items will cause the generation of further interrupts. In the present invention, the interrupt is not enabled when the interrupt is serviced and therefore further work items will not generate interrupts. When the task executes, it processes the work items after which a

10

15

20

25

speculative task is scheduled and added to the task queue.

The task is speculative in the sense that when it is generated there are no work items on the queue to be processed. However it is anticipated that further work items will have been added by the time the speculative task reaches the head of the task queue and is executed.

In accordance with one embodiment of the invention, the method could include the step of continually scheduling speculative tasks (i.e. polling) for processing of work items that may subsequently be received in the system.

However this embodiment, which effectively comprises the use of a single interrupt followed thereafter by polling, is only suitable for systems which have a device which is either never used or is always heavily used during the period of time between system resets.

In a preferred method, when the speculatively scheduled task is executed to process any work items received by the system and it is determined that there are no work items, the interrupt is enabled. Thus when the system is fully utilised, the interrupt mechanism is replaced with a polling mechanism involving a continuous series of speculatively scheduled tasks. However when the system or device utilisation decreases, i.e. when there are no work items when the speculatively scheduled task is processed, then the system reverts to interrupts.

In one alternative embodiment, when device presents a work item to the system, the device operation may stall pending processing of the work item by the scheduled task. This embodiment may provide a benefit in terms of overall system performance in the case where the system throughput is limited by processor utilisation because the interrupt overhead would be eliminated. In a preferred embodiment however, a work item queue is maintained for the interrupting device so that the operation of the device is not halted when the interrupt is disabled.

According to another aspect of the invention there is provided a data processing system comprising: processing means for executing tasks to process work items in the data processing system; and interrupt generating means for generating an interrupt in response to receipt of a work item in the system; wherein the processing means is operable to: service the generated interrupt to schedule a task for later processing of the work item, without re-enabling the interrupt; subsequently execute the task to process the work item; and speculatively schedule a further task for processing of any work items that are subsequently received in the system.

According to a further aspect of the invention there is provided a computer program product comprising a computer usable medium having computer readable program code means embodied in the medium for processing work items in a data

processing system, the program code means comprising: code means for causing the data processing system to service a generated work item interrupt to schedule a task for later processing of the work item, without re-enabling the
5 interrupt; code means for causing the data processing system to subsequently execute the task to process the work item; and code means for causing the data processing system to speculatively schedule a further task for processing of any work items that are subsequently received in the system.

10

The computer program product according to the further aspect of the invention may be implemented in a read-only memory, magnetic tape or diskette.

15
16
17
18
19
20
21
22
23
24
25

The present invention may advantageously be used to reduce overall latency in any interrupt based data processing system where it is expected that there will be a delay between scheduling and execution of tasks.

20
21
22
23
24
25

Brief Description of the Drawings

A preferred embodiment of the present invention will now be described with reference to the accompanying drawings in which:

25

Figure 1 is a schematic representation of a processing apparatus embodying the invention;

Figure 2 is a flow diagram showing the steps involved in a method according to a preferred embodiment of the invention; and

5 Figs 3A, 3B and 3C are schematic representations of the state of the task and work item queues of the preferred embodiment of the invention at different stages of the method of Figure 2.

10 Description of the Preferred Embodiments

In the following the present invention will be described in relation to a storage controller (e.g. storage adapter card) receiving work items from an attached host system. It will be appreciated that this is by way of example only -- the invention may be used in many suitable systems.

15 With reference first to Figure 1, there is shown a processing apparatus comprising a host system 10, which may be a personal computer system, workstation, mainframe or the like, for example an IBM RS/6000 system running the AIX operating system. The host system includes a device driver 20 and a gateway 30 for providing communication between the device driver and an attached storage controller 100. In the 20 present embodiment, the storage controller takes the form of an adapter card which plugs into a bus (for example a PCI bus) within the host system. The storage controller is in

turn attached for communication with a storage subsystem 200 that may comprise an array of data storage disk devices.

The storage controller includes a gateway 102 for
5 communicating with the host system, a microprocessor 104 and DRAM 106 for storing data and metadata, including data being transferred between the host and storage subsystem. Also provided as part of the storage controller is ROM 108 in which is stored firmware that executes on the microprocessor
10 to provide a variety of different services within the controller. The services take the form of a variety of different tasks which are scheduled on a task queue for execution by the microprocessor. One service provided by the firmware is the processing of requests from the host system to read and write data from and to the storage subsystem.
15 Each such data transfer request will generally involve a number of different firmware tasks which during the operation of the controller are placed on the task queue within the DRAM. Disk interface 112 provides the communication interface to the storage subsystem. The controller further includes data flow manager 110 that is adapted to control the data flow within the controller. As will be described below, the data flow manager receives work items from the host system and generates interrupts for
20 servicing by the controller firmware.
25

Next will be described, in general terms, the operation of the processing apparatus of Figure 1. In response to host

application requests for data, the host system device driver generates host transactions (otherwise referred to herein as work items) for processing by the storage controller. The work items are transferred via the gateway to the storage
5 controller and placed on a work item queue in DRAM by the data flow manager. In response to receipt of each work item, a task is caused to be placed on a task queue in DRAM for subsequent execution by the microprocessor. To take the example of a disk write operation, the appropriate task is
10 placed on the task queue for subsequent execution on the processor to DMA the write data from the host system memory into the controller DRAM. It should be noted that tasks other than those involved in processing host work items will also be placed on the task queue for execution by the
15 microprocessor. In the present embodiment, the task queue is managed by the microprocessor as a FIFO queue in a manner well known in the art, with tasks added to the tail of the queue.

20 In prior systems of this type which are interrupt driven, the receipt of a work item at the controller will result in the generation of an interrupt to the microprocessor. As discussed previously, such a system provides low latency and low overhead provided that the rate
25 of interrupts is low. If however, the number of work items from the host is significantly increased, the interrupt overhead will be increased to a level which may impact the overall operation of the controller microprocessor.

With reference next to Figure 2 and Figures 3A to 3C, there will be described a process, operable on the storage controller of Figure 1, for processing work items from the host system in a manner that combines the best attributes of the polling and interrupt methods to service the host work items with low latency at low utilisation and by polling for low overhead at high utilisation.

In accordance with the present embodiment, the process starts at step 300 where the controller is initialised. At step 310, the host work item interrupt is enabled ready to allow the generation of a host work item interrupt in response to receipt of a work item at the controller. At step 320, a host work item is received by the controller data flow manager and placed on the host work item queue in DRAM. As the interrupt is enabled, the data flow manager generates an interrupt to the microprocessor which causes the microprocessor to halt its operation and to load the appropriate interrupt handler code from ROM. At step 330, the microprocessor disables the interrupt thereby preventing the generation of any host interrupts pending processing of the work item. At step 340, the interrupt handler schedules a task for processing the work item by placing the task on a FIFO task queue within the DRAM. The interrupt handler then completes and the processor returns to the operation it was carrying out prior to the interruption. In contrast with prior systems, the interrupt is not enabled at this point;

this means that if and when a further work item is received on the work item queue, an interrupt is not generated. The status of the work item and task queues at this stage of the process can be seen in Figure 3A. The work item queue
5 includes one item I1 and the task queue includes a number of tasks including T1 at the tail of the task queue.

Unless it is subject to other interruptions from other devices, the microprocessor continues by processing the
10 tasks on the queue and eventually task T1 reaches the head of the queue. During the period between when task T1 was placed on the queue and when T1 reached the head of the queue, three further work items have been queued by the data flow manager. As the interrupt was disabled, no interrupts were generated in response to receipt of these work items.
15 The status of the work item and task queues at this stage of the process can be seen in Figure 3B in which the work item queue includes three items I1 to I3 and the task queue includes a number of tasks including task T1 at the head of the queue.
20

At step 350 the scheduled task T1 is loaded and executed by the microprocessor to process the work item I1
25 (step 360). At step 370, the processor checks for further work items on the work item queue and on finding item I2, processes that as well. The processor continues to check for and process further work items until the queue is empty.

In accordance with the embodiment of the invention, on finding the queue to be empty, the processor, at step 380, causes a further task for processing host work items to be scheduled by placing it on the task queue. The interrupt is
5 not enabled/unmasked. The status of the work item and task queues at this stage of the process is as indicated in Figure 3C in which the work item queue is empty and the task queue includes a number of tasks including the speculatively scheduled task at the tail of the queue.

10

This further task is speculative in nature in the sense that the processor is anticipating that further work items will appear on the queue in the meantime and therefore when the speculative task reaches the head of the queue, it will have work items to process. Because the interrupt is not enabled when the speculative task was scheduled, none of the work items which are subsequently added to the queue will generate interrupts, thus avoiding interrupt overhead.

15

Again assuming that no further interrupts are received, the microprocessor continues by processing the tasks on the queue until task T2 reaches the head of the queue. At step 390, the microprocessor executes task T2 which determines
20 (step 400) whether there are any work items on the work item queue. If yes, the task processes the queued work items at step 410. On a determination that there are no further work items, the processor causes a further speculative task to be scheduled (step 380) and placed on the task queue. Although
25

this scheduling of a speculative task involves a certain amount of overhead, this is worthwhile in view of the likelihood that the work items will continue to be placed on the queue.

5

If at step 400 it is determined that there are no host work items on the queue, then the process moves to step 310 and the interrupt is enabled/unmasked. Thus when no further work items have been added to the queue in the time between the task T2 is added to the task queue and when it reaches the head of the queue, the host and/or controller workload is low and the process reverts to an interrupt driven process for the host work item interrupts.

10

15

20

25

Thus it will be noted that when the controller is relatively busy (i.e. when there are many different tasks on the task queue), the speculatively scheduled task will be delayed by other tasks which are currently in progress and thus the length of the delay will increase according to how busy the controller is. During the period the speculative task is delayed, work items from the host will be added to the queue. Because the interrupt is disabled/masked, none of these work items results in the generation of an interrupt and therefore the interrupt overhead which would occur in prior systems is avoided. When the speculative task reaches the head of the queue all the pending work items are processed.

When both the arrival rate of host work items and the controller utilisation are low, the pattern of operation is that an arriving host work item generates an interrupt which schedules a task to service the work item. The task services the work item and a speculative task is scheduled. The speculatively scheduled task is executed a relatively short time after the interrupt scheduled task because there are few other outstanding tasks to execute on the controller between the interrupt scheduled task and the speculatively scheduled task. The short delay between the interrupt scheduled task and the speculatively scheduled task and the low arrival rate of host work items means that the speculatively scheduled task does not find host work items waiting when it polls for them so the speculative polling process is terminated and the interrupt mechanism is restored after only one speculatively scheduled task. Thus the host system is serviced with the low latency of the interrupt method with some additional overhead caused by the speculatively scheduled tasks. This additional overhead does not impact the overall operation of the controller because the system is relatively idle and the controller is under-utilised.

As the arrival rate of host work items increases and/or the controller utilisation increases (due for instance to an increased arrival rate of work items from a different source), there is more opportunity for host work items to arrive between the interrupt scheduled task and the

speculatively scheduled task (or between a speculatively scheduled task which found host work items and the speculatively scheduled task which was scheduled as a result) so the speculatively scheduled task will more frequently find host work items waiting and the speculative polling process will go on for longer before the interrupt mechanism is restored. Thus the number of interrupts avoided and the size of the 'batches' of work items processed by each task increases both with the arrival rate of host work items and with the utilisation of the controller.

Although on the face of it, it might appear that by use of the technique described above, the delays in processing the work items might increase the latency of the system at high utilisations over that of the conventional interrupt method, this is not in fact the case. The reduced overhead of this technique more than compensates by keeping the queue depth down. In the limit, the present invention allows the system to operate with relatively low latency at throughputs beyond the limit where the conventional interrupt method would have caused the queue depth and latency to tend to infinity.

While an embodiment of the invention have been described in detail above, it will be apparent to those skilled in the art that many variations and modifications can be made without departing from the scope of the

invention. In particular, although in the described embodiment, the queued work items were host transactions, it will be appreciated that the use of the invention in a storage controller is not limited to such operations. For
5 example, the technique of the present invention may be used with other interrupt-generating events such as DMA completions and messages received from an attached storage controller. Multiple queues would be maintained in the event that the technique of the present invention is used in
10 relation to multiple events.

Furthermore, although the invention has been described in relation to a storage controller, it will be apparent that the invention may be used in any data processing system in which interrupts are generated in response to the receipt of work items from a device.
15

In addition, although the work items are described as being managed in a queue, in alternative embodiments, the work items may be arranged as an array where a work item exists in a fixed location and is marked as requiring servicing by the hardware setting a bit in the array.
20

CLAIMS

1. A method for processing work items in a data processing system comprising:

5

generating an interrupt in response to receipt of a work item in the system;

10

servicing the generated interrupt to schedule a task for later processing of the work item, without re-enabling the interrupt;

15 subsequently executing the task to process the work item; and

20

speculatively scheduling a further task for processing of any work items that are subsequently received in the system.

25

2. A method as claimed in claim 1 comprising the further steps of:

executing the speculatively scheduled task to process any work items received by the system;

on a determination that there are no work items to be processed, enabling the interrupt; and

on a determination that there are work items to process, speculatively scheduling a further task, without re-enabling the interrupt.

5 3. A method as claimed in claim 1 wherein the work items are managed on a queue.

10 4. A method as claimed in claim 1 where in the event that further work items are received after the task is scheduled and prior to execution of the task, the step of executing

speculatively schedule a further task for processing of any work items that are subsequently received in the system.

6. A data processing system as claimed in claim 5, the
5 processing means being operable on a determination that there are work items to be processed to execute the speculatively scheduled task to process the work items and to schedule a further speculative task; and operable on a determination that there are no work items to be processed
10 to enable the interrupt.

7. A data processing system as claimed in claim 5 further including memory for storing the received work items on a queue.

8. A data processing system as claimed in claim 5 where in the event that further work items are received after the task is scheduled and prior to execution of the task, the processing means is operable to execute the task to process all the work items.

9. A data processing system as claimed in 5 wherein the interrupt generating means and processing means are embodied in a data storage controller and the work items comprise
25 data transfer requests from an attached host system.

10. A computer program product comprising a computer usable medium having computer readable program code means embodied

in the medium for processing work items in a data processing system, the program code means comprising:

5 code means for causing the data processing system to service a generated work item interrupt to schedule a task for later processing of the work item, without re-enabling the interrupt;

10 code means for causing the data processing system to subsequently execute the task to process the work item; and

15 code means for causing the data processing system to speculatively schedule a further task for processing of any work items that are subsequently received in the system.

11. A computer program product as claimed in claim 10, the computer readable program code means further comprising:

20 code means for causing the data processing system to execute the speculatively scheduled task to process any work items; and

25 code means for causing the data processing system to enable the interrupt on a determination that there are no work items for processing.

ABSTRACT OF THE DISCLOSURE

Described is a method for processing work items in a data processing system. An interrupt is generated in response to receipt of a work item on a queue and the generated interrupt is serviced to schedule a task by placing the task on a task queue for later processing of the queued work item. The interrupt is not enabled at this point and therefore the receipt of further work items will not result in the generation of interrupts. When the scheduled task reaches the head of the queue, the task is executed to process the queued work item (and all other work items that have been added to the queue since the task was scheduled). When all the work items have been processed, a further task is speculatively scheduled for processing of any work items that are subsequently placed on the work item queue. When the speculatively scheduled task reaches the head of the queue, it is executed to process queued work items; if any work items were found, a further task is speculatively scheduled, otherwise the interrupt is enabled.

5

10

15
16
17
18
19
20

20

^{1/3}
H.E. Butterworth
UK 999-027

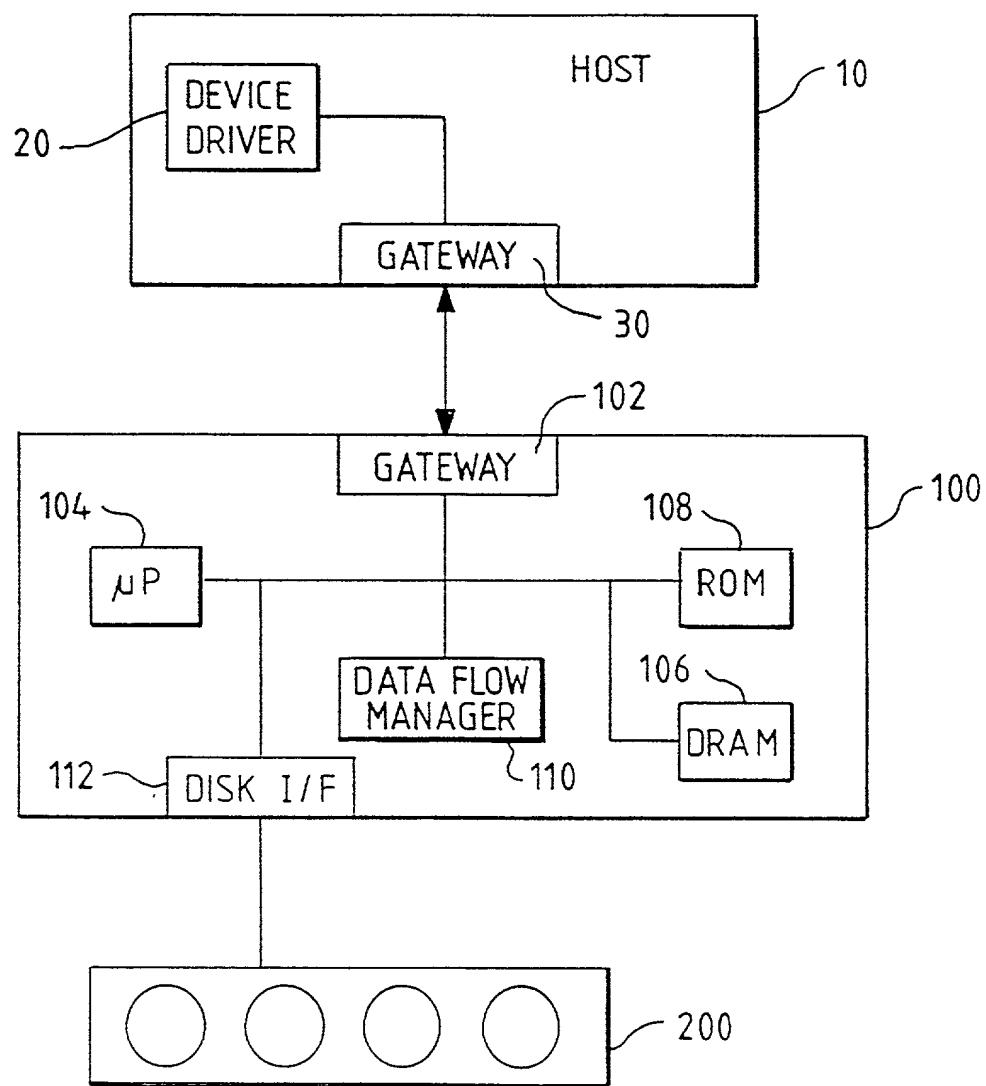


FIG. 1

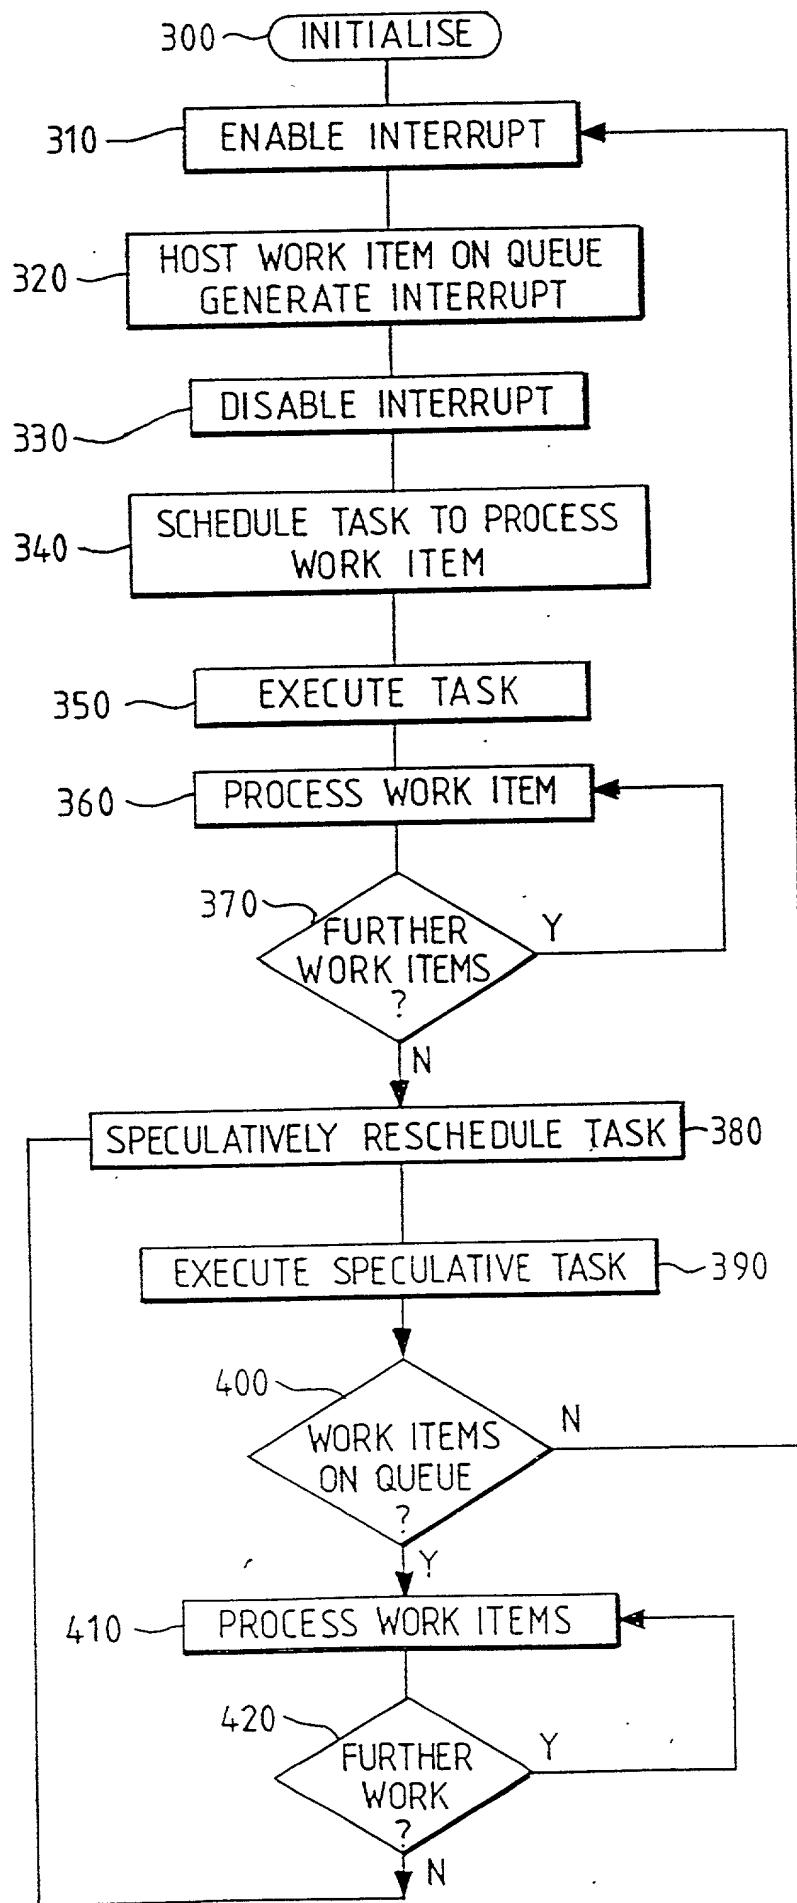
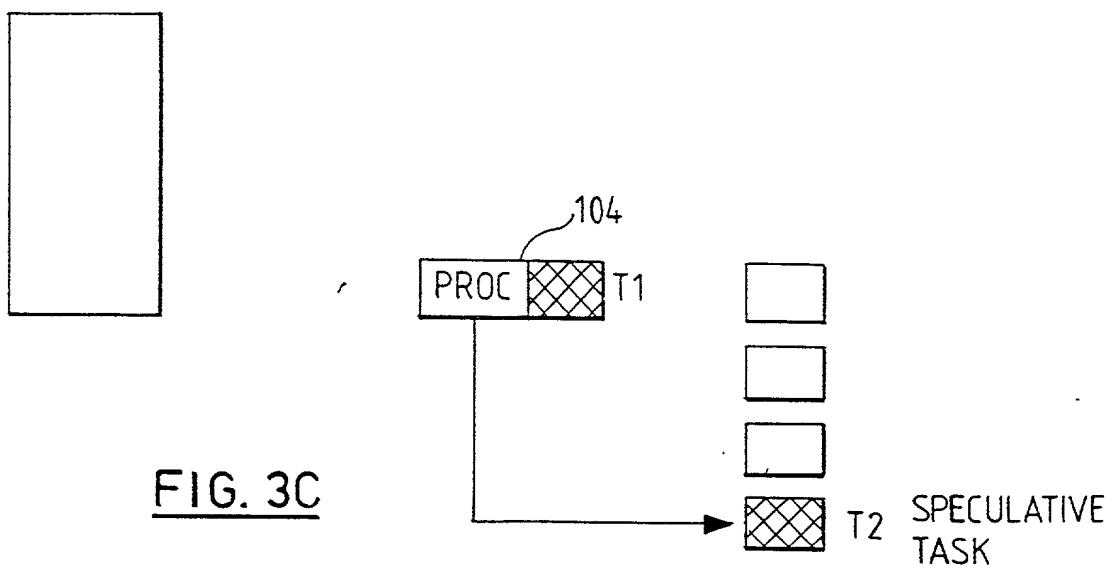
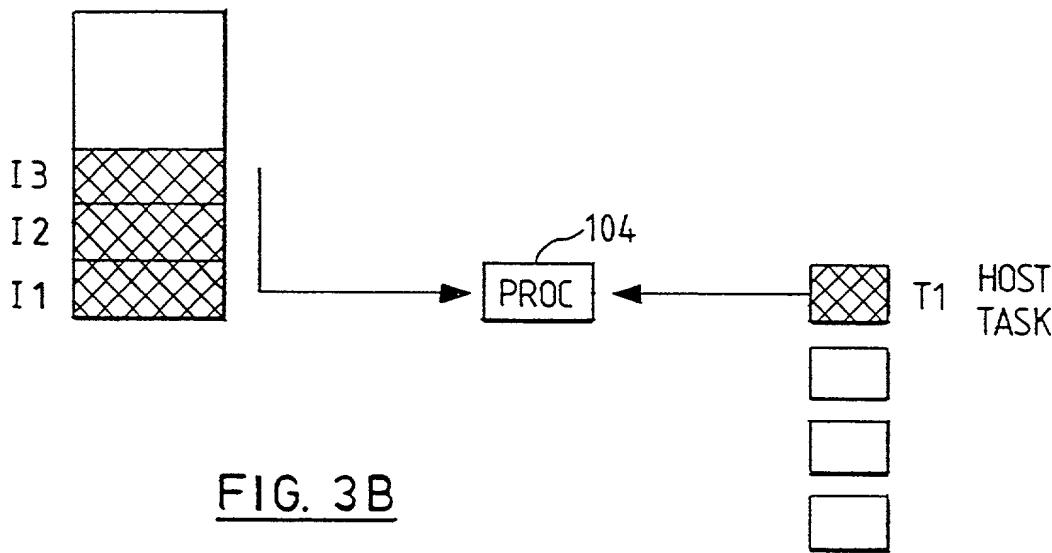
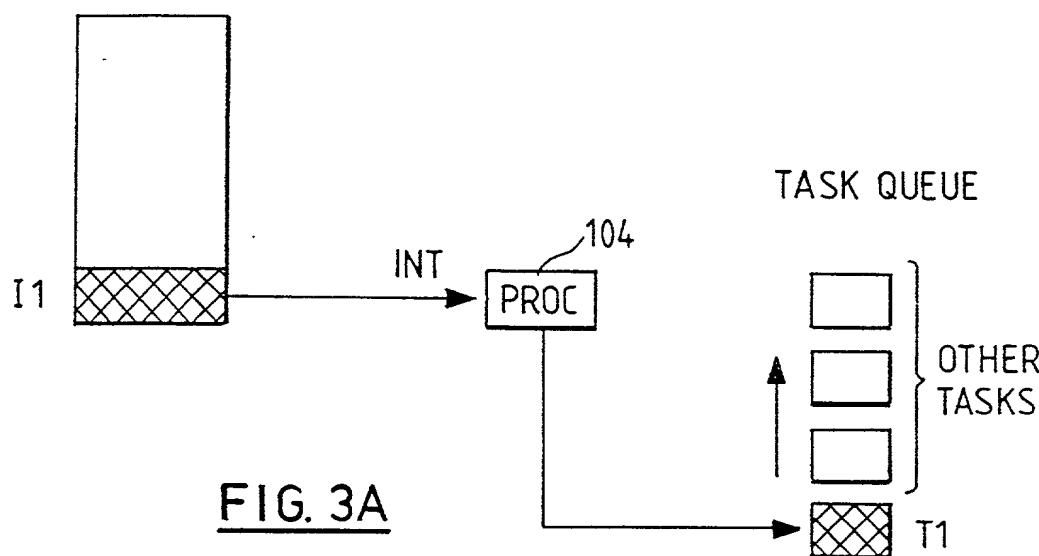


FIG. 2

HOST WORK ITEMS



DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name; I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

DATA PROCESSING SYSTEMS AND METHOD FOR PROCESSING WORK ITEMS IN SUCH SYSTEMS

the specification of which (check one)

is attached hereto.

was filed on _____
as Application Serial No. _____
and was amended on _____

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, CFR 1.56.

I hereby claim foreign priority benefits under Title 35, USC 119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s):

9906501.3	United Kingdom	23/03/99	<input checked="" type="checkbox"/> Yes <input type="checkbox"/> No
Number	Country	Filing Date	Priority Claimed
		Day/Month/Year	

I hereby claim the benefit under Title 35, USC 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, USC 112, I acknowledge the duty to disclose information material to the patentability of this application as defined in Title 37, CFR 1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

Prior U.S. Applications:

Serial No.	Filing Date	Status
------------	-------------	--------

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Manny W. Schechter, Reg. No. 31,722; William B. Porter, Reg. No. 33,135; Douglas W. Cameron, Reg. No. 31,596; Kevin M. Jordan, Reg. No. 40,277; Stephen C. Kaufman, Reg. No. 29,551; Richard M. Ludwin, Reg. No. 33,010; Daniel P. Morris, Reg. No. 32,053; Louis J. Percello, Reg. No. 33,206; Jay P. Sbrollini, Reg. No. 36,266; Stephen S. Strunck, Reg. No. 28,672; Robert P. Tassinari, Reg. No. 36,030; Robert M. Trepp, Reg. No. 25,933; Marc D. Schechter, Reg. No. 28,989; Louis P. Herzberg, Reg. No. 41,500.

Send all correspondence to:

IBM Corp, IP Law Dept,
T J Watson Research Center, PO Box 218,
Yorktown Heights, New York 10598.
Telephone: (914)945-

FULL NAME OF SOLE OR FIRST INVENTOR: **Henry Esmond Butterworth**

INVENTOR'S SIGNATURE:

 DATE: 21 May 1999

RESIDENCE:

23 Maytree Road, Chandlers Ford, Hampshire SO53 5RT,
United Kingdom

CITIZENSHIP:

British

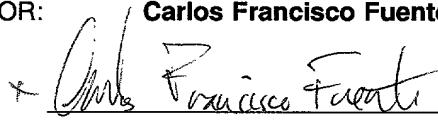
POST OFFICE ADDRESS:

As above

FULL NAME OF SECOND INVENTOR:

Carlos Francisco Fuente

INVENTOR'S SIGNATURE:

 DATE: 21 May 1999

RESIDENCE:

22 Delamere Road, Southsea, Portsmouth, Hampshire PO4
0JA, United Kingdom

CITIZENSHIP:

British

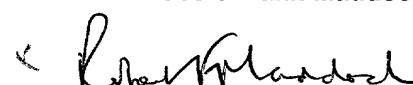
POST OFFICE ADDRESS:

As above

FULL NAME OF THIRD INVENTOR:

Robert Frank Maddock

INVENTOR'S SIGNATURE:

 DATE: 21 May 1999

RESIDENCE:

34 Welbeck Avenue, Southampton, Hampshire SO17 1SU,
United Kingdom

CITIZENSHIP:

British

POST OFFICE ADDRESS:

As above

Attorney Docket No. UK999-027

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICANT(S): Henry Esmond Butterworth
Carlos Francisco Fuente
Robert Frank Maddock

SERIAL NO.: Unassigned

FILED: Concurrently Herewith

FOR: DATA PROCESSING SYSTEMS AND METHOD FOR PROCESSING WORK ITEMS IN SUCH SYSTEMS

ASSOCIATE POWER OF ATTORNEY

Please recognize JOSEPH B. RYAN, Reg. No. 37,922; KEVIN M. MASON, Reg. No. 36,597; and WILLIAM E. LEWIS, Reg. No. 39,274; each of them of RYAN & MASON, L.L.P., 90 Forest Avenue, Locust Valley, New York 11560 as associate attorneys in the above-mentioned application, with full power to prosecute said application, to make alterations and amendments therein, and to transact all business in the Patent and Trademark Office connected therewith.

Telephone calls should be made to William E. Lewis by dialing (516) 759-2946.

All written communications are to be sent to William E. Lewis, Esq., Ryan & Mason, L.L.P., 90 Forest Avenue, Locust Valley, New York 11560.

Dated: 9/22/99

✓ Barthole

Manny W. Schecter
Registration No. 31,722
Stephen C. Kaufman
Registration No. 29,551
Jay P. Sbrollini
Registration No. 36,266
Attorney for Applicant(s)

International Business Machines
Corporation
T.J. Watson Research Center
Route 134 and Kitchawan Road
Yorktown Heights, New York 10598